Seizure Detection and Prediction Through Clustering and Temporal Analysis of Micro Electrocorticographic Data

Yilin Song [1], Jonathan Viventi [2], and Yao Wang [1]

[1]Department of Electrical and Computer Engineering, New York University, NY, USA
[2]Department of Biomedical Engineering, Duke University, Durham, NC, USA

***Abstract*—We have developed flexible, active, multiplexed recording devices for high resolution recording over large, clinically relevant areas in the brain. While this technology has enabled a much higher-resolution view of the electrical activity of the brain, the analytical methods to process, categorize and respond to the huge volumes of data produced by these devices have not yet been developed. In our previous work, we applied advanced video analysis techniques to segment electrographic spikes, extracted features from the identified segments, and then used clustering methods (particularly Dirichlet Process Mixture models) to group similar spatiotemporal spike patterns. From this analysis, we were able to identify common spike motion patterns. In this paper, we explored the possibility of detecting and predicting seizures in this dataset using the Hidden Markov Model (HMM) to characterize the temporal dynamics of spike cluster labels. HMM and other supervised learning methods are united under the same framework to perform seizure detection and prediction. These methods have been applied to in-vivo feline seizure recordings and yielded promising results.**

I. INTRODUCTION

Currently, many existing neurological data analyses rely on manual inspection. With new high-density, micro-electrocorticographic (μECoG) electrode arrays that provide dramatically enhanced spatial resolution [1], the volume of data produced by arrays with hundreds or thousands of electrodes is too large for manual review. Further, manual inspection can miss subtle features that automated machine learning techniques can detect. There is an urgent need for efficient and sensitive automated methods that can analyze the large volumes of data produced by next generation neurologic devices.

One major application of high resolution μECoG is to record brain signals in patients with epilepsy. Preliminary analysis of our high resolution data have discovered repetitive spatiotemporal microscale spike patterns that appear to initiate and terminate seizures. These patterns are likely not observed by standard electrodes, as their spatial scale is smaller than the current electrode spacing (approximately 1cm). Understanding the relationships between these spike patterns is a key step towards developing a deeper understanding of the microscale phenomena that generate seizures, and ultimately enable us to develop better therapies to treat epilepsy.

In [2], we compared many video analysis and machine learning techniques to recognize the spatiotemporal patterns of electrographic spikes. We have developed a three-layer, pipelined, framework for robust spike segmentation and pattern recognition. In the first layer, we extracted spike segments from raw signal recording via spatial temporal filtering and region growing[3]. In the second layer, we extracted features from these segments. Finally, in the third layer we used unsupervised learning to group recurrent patterns into clusters. The three layers were reflected in Fig. 2 as the first three layers in the diagram. In this work, we investigate how to use the temporal variation of the spike pattern label to detect and predict seizure, as indicated in the fourth layer in Fig. 2.

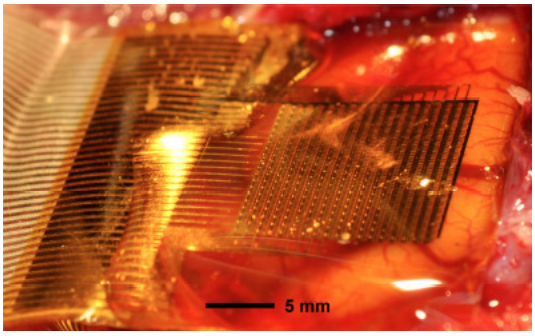

Fig. 1: Photograph of a 360 channel, high density neural electrode array consisting 18×20 electrodes used in a feline model of epilepsy. The electrode array was placed on the surface of visual cortex. The electrode size and spacing was 300 μm x 300 μm and 500 μm, respectively.

Most current seizure detection and prediction methods fall into three categories: 1) Examination of the EEG signals in the frequency domain using wavelet transform. [4], [5] has demonstrate that wavelet transform is a good feature representation to differentiate abnormal event from normal brain activities. [6] has extended to seizure detection with additional RBF neural network structure. [7], [8] has demonstrated SVM would have additional prediction power operating on wavelet features. 2) Analysis of the linear or nonlinear temporal evolution of the EEG signals to find a governing rule. [10] used Kalman filter for modeling EEG signals. [11] has shown SVM operating on autoregressive coefficients of EEG signals would have seizure

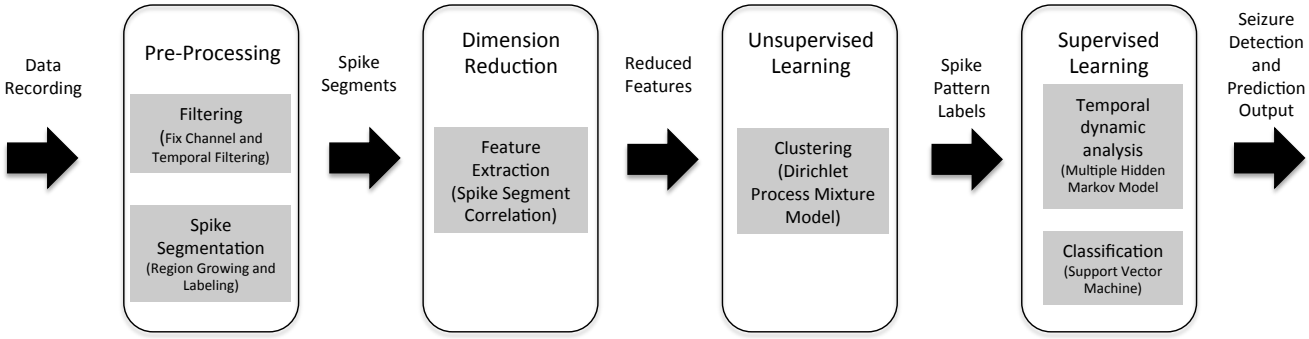

Fig. 2: Block diagram showing 4-stage seizure detection and prediction algorithm. In data preprocessing, we spatially interpolates missing channels and filters each channel with a band pass filter

prediction power. There have been several works using hidden markov model (HMM) [14] for seizure detection. Some of these works are based on single channel signal [15]. [16] has shown that with additional SVM structure before HMM would eliminate noise so that the frame work would have prediction power. The work in [17] [18] consider multiple channel signals for seizure detection and are most closely related to our proposed research. Most of these prior work use a single HMM, with the hidden states representing the stage of seizure to be identified. 3) Use eigenvectors of space-delay correlation and covariance matrices as feature representation followed by classifiers [12], [13]. To best of the authors' knowledge no research has been published on an integrated and simultaneous application of spike segmentation, clustering and temporal dynamic analysis for seizure prediction.

If spikes occur less frequently in future datasets from more accurate models of epilepsy, the length of the fixed length segment partition can be increased, which could allow our method to predict seizures farther in advance. In our approach, we investigate seizure detection and prediction through classifying a spike into one of the three classes (non-seizure, pre-seizure, and seizure). A spike that is classified into the pre-seizure stage indicates the on-set of a new seizure. We first classify this spike into one of the pre-determined spike clusters based on features extracted from this spike signal. Seizure stage classification is based on how the the spike cluster label changes in time, over a sliding time window surrounding the current spike time. This is accomplished through HMM modeling of the dynamics of spike labels in different seizure stage. If there were inactive period between two spikes, this period will be partitioned into fixed length segments (with length equal to the average spike segment length), each given a special label that corresponds to its inactive status. If spikes occur less frequently in future datasets from more accurate models of epilepsy, the length of the fixed length segment partition can be increased, which could allow our method to predict seizures farther in advance.

The rest of this paper is organized as follows: In Section II, we review our prior work on spike segmentation and clustering. In Section III we present seizure detection and prediction algorithm using the clustering label. Section IV concludes the paper with a discussion of the results.

II. Spike Segmentation and Clustering

A. Spatial-Temporal Segmentation

In our acquired dataset, the multi-channel μECoG signal has a high spatial and temporal correlation, and therefore each spike could be treated as a spatio-temporally connected set of voxels that have high intensity values. We applied 3D region growing to get spike segments. Region growing first randomly select pixels (called seeds) that have intensity values above a preset threshold P_t. Each seed is used as an initial detected region. The algorithm then iteratively examines whether any immediate neighbors of the boundary of previously detected regions, S_n, also have high intensity values, with a threshold that was determined based on the mean, η_n, and standard deviation, σ_n of the pixels inside the region for each iteration n. Specifically, a neighboring pixel with intensity P was included in S_n if $P >= \eta_n - \alpha\sigma_n$. This process continues until no more pixels could be included in the region S_n. We applied the region growing algorithm to the data captured from an acute in vivo feline model of seizures, using the parameters $P_t = 0.8mV$, $\alpha = 0.8$. We further ignores regions last no longer than $40ms$ after the algorithm converges. Figure 3 shows two detected spike segments from one dataset.

B. Feature Extraction

Clustering the spike patterns based on raw data presents several challenges. First, spike segments will have varying time durations, leading to different dimensions of the feature

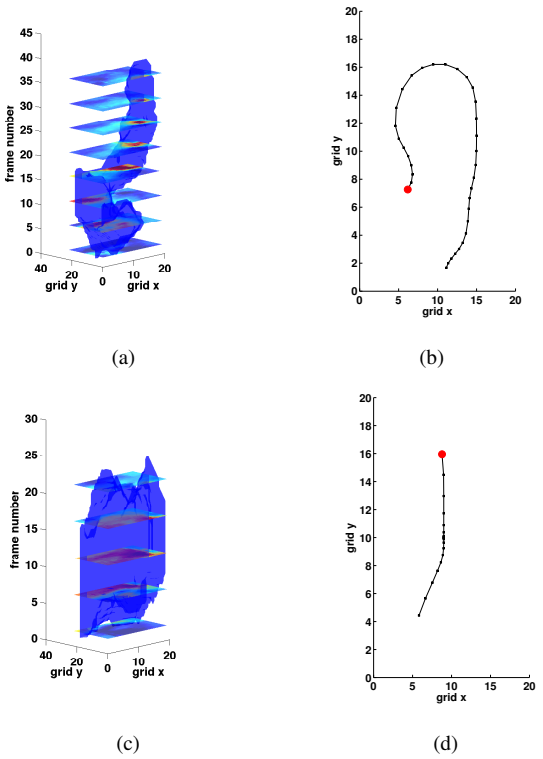

Fig. 3: Examples of spike segmentation. Subfigure a,b are for a spike with a spiral motion: Subfigure a shows the segmented volume in blue overlaid with the μECoG signals captured at different times, with vertical axis (time axis) corresponding to frame number, with the other two axis represent the spatial arrangement of μECoG electrodes; Subfigure b shows the trajectory of the centroid of the segmented region in successive frames, with the red dot indicating the beginning position; Subfigure c,d are for a spike with a vertical motion.

vectors. Second, the dimensionality of such feature vectors will be very large.

To circumvent the above challenges, we used the correlation coefficients between each spike segment with all other spike segments in a training dataset to form a feature vector. Assuming that the signals corresponds to two spike segments have duration n and m with $n \geq m$. We denoted their corresponding video signals by $X \in R^{N_1 N_2 n}, Y \in R^{N_1 N_2 m}$, where the intensity values at voxels not belonging to the spike segment are set to zero. To compute the correlation between X and Y, we found a length-m segment in X that has the highest correlation coefficient with Y, and use this maximum as the correlation between X and Y.

The raw signal correlation map has a dimension of $N * N$, if N is the total number of segments in the training set. Each column in this correlation map served as the feature vector for one spike segment. The feature vector consists of the correlation of this spike segment with all other segments. The motivation for using the correlation as the feature vector for clustering was that similar segments should have similar correlations with respect to other segments. Correlation features may not depict the spatio-temporal characteristics of the spike segment directly, but have proven to be more effective than other features we explored [2], for unsupervised clustering of spike patterns.

C. Unsupervised Clustering of Individual Spike Pattern

In order to identify all typical spike motion patterns, we perform unsupervised clustering of all spike segments in a training set. Among the various clustering methods that we explored in [2], the Dirichlet Process Mixture Model (DPM) renders more consistent clusters under two of our evaluation metrics. DPM is a special mixture model, where the mixing coefficient is a random variable that follows a certain distribution. We used Mean-Field Variational Inference[19] to learn the DPM parameters that best described a given set of samples. With this algorithm, only an upper bound needs to be provided to set the maximum possible number of mixtures, it automatically determines the mixture number that leads to the best fit (in terms of the likelihood) of the training data and the model found. The DPM-based clustering is particularly effective when the training data tend to form clusters that vary greatly in size (i.e. some clusters contain many more samples than other clusters), which is the case with our data. We used a training set to find the cluster number and the probability distribution of each cluster. For seizure detection and prediction, for each new spike segment, we first derive its feature vector by finding its correlation with all spike segments in the training set, and then assign it to the cluster with the maximum likelihood.

III. Seizure Detection and Prediction

Once a spike is detected and clustered into one of the predetermined clusters, we would like to classify it into one of three spike categories: interictal, pre-ictal, ictal (i.e seizure) based on the temporal variation pattern of the clustering label. For neural datasets, a significant portion of the data recording is low amplitude signal, for which we assign a unique label 0. For resting period d longer than 100 ms, the number of zero labels is $ceil(d/100)$, since the average duration for spike segments is 100 ms. To train the classifier, we manually label the ictal stage of each spike segment in training sequences, and assign spikes that occur 8 seconds before ictal stage to be pre-ictal stage spikes.

A. Hidden Markov Model

Hidden Markov Model is quite powerful for temporal data analysis, especially when the observations at successive times can be effectively described as discrete variables. The hidden state q_t is assumed to consist of N possible values. The probability of $q_t = i$ is conditionally independent of all other previous variables given q_{t-1}. The observation O_t at time t is assumed to vary among K possible values, and is conditionally independent of all other variables given q_t. Parameters of HMM are defined as following:

$$A_{i,j} = P(q_t = S_i | q_{t-1} = S_j)$$
$$B_j(k) = P(O_t = k | q_t = S_j)$$
$$\pi_i = P(q_1 = S_i)$$

where A_{ij} are called transition probabilities, $B_j(k)$ the emission probabilities of state i, and π_i the stationary probability of state i.

B. Using a Single Hidden Markov Models

We have explored two different ways of using HMM for seizure stage classification. In this single HMM approach, hidden states in the HMM correspond to all possible seizure stages (i.e., interictal, pre-ictal, and ictal). Given a sequence of T observations (observation at time t is the cluster label of the spike at time t), the most likely state transition path then tells the classification results for all T spikes. During the training stage, we collect observations and their corresponding state labels to derive the transition and emission probabilities for HMM. Inferring the hidden state label in this case corresponds to finding the Viterbi path. We found that prediction accuracy begins to decrease as the observation sequence length reached beyond a certain point. Therefore instead of using all past observations up to time t, we use only a fixed number of past observations. Through trial and error, we found that using 20 observations yielded the best result. We have experimented with two ways for assigning the initial state of HMM, 1) using the stationary probability, 2) using the state assigned to the first spike segment when applying HMM classifier to the previous overlapping observation sequence.

C. Using Multiple Hidden Markov Models plus a second stage classifier (HMM+SVM)

In this approach, we built a separate HMM for each individual seizure stage. $\lambda_i = (A_i, B_i, \pi_i)$ denotes the parameters for the i-th stage. During the training stage, the goal is to adjust model parameters $\lambda_i = (A, B, \pi)$ to maximize $P(O|\lambda_i)$ using observation sequences in the training data for stage i. Model parameters are updated iteratively through Baum-Welch algorithm [14]. Since we have no prior knowledge of how many states to choose, the number of hidden states n is chosen to maximize $P(O|\lambda_i)$. We found that four hidden states work well for all three different HMMs.

We have found that using the HMM trained separately for each seizure stage cannot accurately classify a spike during the transition period from non-seizure to seizure (or vice versa). We suspect that this is because, during the transition period, the observation sequence contains spikes from both the non-seizure period and the seizure period. If we record the actual likelihood values for all three seizure stages in time, it is likely that, during the transition period, the likelihood for the non-seizure stage decreases, and the likelihood for the seizure stage increases. To exploit the likelihood variation pattern over time, we adopted a sliding window approach and find the likelihood vector (consisting of the likelihood values of the three seizure stages) over each sliding window. As shown in Fig 4, to classify the observation O_t as indicated by the red double dashed vertical lines on the top, we collect a total of $2n + 1$ observations. For each sliding window containing n observations, we derive the log likelihood that it belongs to each of the three seizure stages using the 3 trained HMM models. Because there are a total of $n + 1$ sliding windows over $2n+1$ observations, we form a new $3n+3$ feature vector G_t that contains the log likelihood values for the three seizure stages of each over the $n + 1$ sliding windows. We apply a trained classifier to this feature vector to determine the seizure stage of the spike O_t. For this second stage classifier, we use the multi-class support vector machine with linear kernel and radial basis function (RBF) kernel.

Dataset	sampling frequency	recording length	labeled seizures	spikes segments
cat 1	277.778 Hz	53 min 41 sec	7	1685
cat 2	925.925 Hz	32 min 20 sec	27	3706
cat 2	925.925 Hz	26 min 10 sec	27	2380

TABLE I: Statistics of the datasets collected from two different animals. Dataset 2 and Dataset 3 are from the same animal, with different implant position.

D. Dataset, Detection and Prediction Result

We analyzed micro-electrocorticographic (μECoG) data from an acute in vivo feline model of epilepsy. Adult cats were anesthetized with a continuous infusion (3 $\sim$ 10 mg/kg/hr) of intravenous thiopental. A craniotomy and durotomy were performed to expose a 2 x 3 cm region of cortex. The high resolution electrode array was then placed on the surface of the brain over primary visual cortex, localized by electrophysiological recordings of visual evoked potentials. Picrotoxin, a GABA-A receptor antagonist that blocks inhibition, was topically applied adjacent to the anterior-medial corner of the electrode array in an amount sufficient to induce abnormal electrical spikes and seizures from the covered region[1].

The active electrode array placed on the cortex was used to record data from 360 independent channels arranged in 20 columns and 18 rows, spaced $500\mu m$ apart. Each electrode contact was composed of a $300\mu m \times 300\mu m$ square of platinum. Two high-performance, flexible silicon transistors for each electrode buffered and multiplexed the recorded signals [1]. The total array size was $10mm \times 9mm$.

Because of the variation of implant position and time effect, the predictor we built is dataset specific. But to avoid overfitting, we use leave one out cross validation. Namely we use observation sequences extracted during one seizure period, pre-seizure period (8 sec) before this seizure, and the non-seizure period leading up to this seizure as our testing data, use the rest of the data for training. We repeat this approach for each seizure.

We compared the performances of both the single HMM classifier and the HMM+SVM classifier (Fig. 4) for determining the seizure stage of a spike. The spike classification performance using the single HMM classifier is given in Table II , results obtained using the HMM+SVM approach are summarized in Table III. Here the accuracy is defined as the

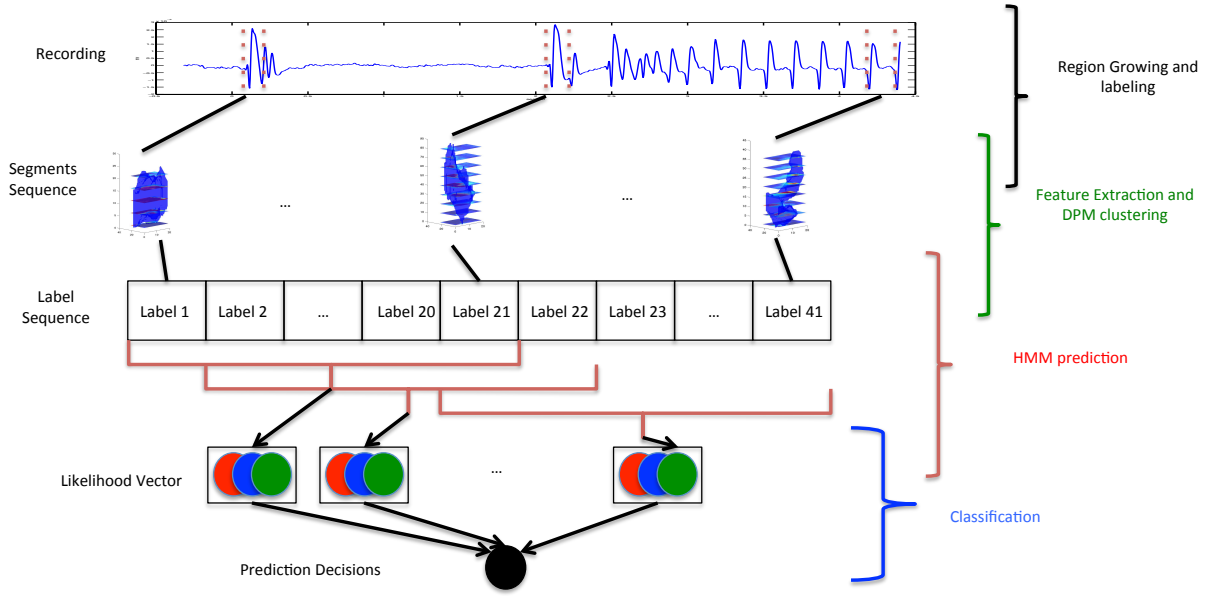

Fig. 4: Illustration of the HMM+SVM seizure prediction scheme. Input of the diagram is the raw μECoG data. Black bracket indicates the process of segmenting the raw data into spike segments. Green bracket generalizes the procedure of extracting features for each spike and classifying it into one of predetermined spike clusters using pre-trained DPM model. Red bracket generalizes the process of multiple HMM modeling for each overlapping sequence of 21 spike cluster labels. Three different colors in each box represent the likelihood of the sequence belonging to one out of three HMMs. Blue bracket represents SVM classification of the likelihood vector computed from a total of 21 overlapping sequences centered around the current spike. The output of the diagram is the classified seizure stage for each spike. Seizure onset is predicted when the current and previous 4 spikes are all classified as either pre-ictal or ictal spikes.

percentage of spikes that are correctly identified into the three seizure stages as compared to the manually labeled ground truth. We can see that the HMM+SVM approach achieves significantly higher accuracy than the single HMM approach.

Tables II and III considered the spike classification accuracy. Even though the HMM+SVM classifier was able to achieve very high spike classification on average, the classification accuracy for the pre-ictal spike is much lower, because pre-ictal spikes occur much less frequently than inter-ictal and ictal spikes. Furthermore, we cannot simply consider any spike classified as pre-ictal as the onset of a new seizure, as that will yield very noisy and inconsistent prediction. Instead, we predict the on-set of a new seizure if the current spike and 4 previous spikes are classified to either pre-seizure or seizure spike. If the seizure actually happened after the predicted onset within 8 sec, we consider the prediction accurate. Table IV summarizes the seizure prediction accuracy using the HMM+SVM approach. We further measure the delay between the predicted on-set and the actual on-set, and report the mean delay in Table IV. Figure 5 shows the distribution of the prediction delay.

Delay	Stationary initial	Recursive initial
no delay	0.765	0.770
10 segments ($\approx$ 1 sec delay)	0.779	0.784

TABLE II: Spike classification accuracy for dataset 2 using the single HMM classifier. The observation sequence length is fixed at 21 spike segments, spike segment to be predicted is either located in the middle of the sequence(delay approximately 1 sec) or in the end (no delay). "Stationary initial" means that we used the stationary state distributions to set the initial state of each observation sequence. "Recursive initial" means that we use the state assigned to the first spike segment by the classifier for the previous observation sequence.

Observation length	svm-rbf	svm-linear
21 segments ($\approx$ 1 sec delay)	0.920	0.846
41 segments ($\approx$ 2 sec delay)	0.940	0.852

TABLE III: Spike classification accuracy for dataset 2 using the HMM+SVM classifier. Spike segment to be predicted is located in the middle of the observation

IV. CONCLUSION

μECoG has tremendous potential for many research and clinical applications. In this work, we have combined advanced video analysis and machine learning algorithms to analyze these challenging datasets in novel ways. From our previous work [2], we have developed efficient methods for identifying and localizing the spatial and temporal extent of inter-ictal and ictal spikes and detecting the spike wavefront through region growing. For those identified spikes we derive correlation features relying on their relationship with one another. DPM has rendered consistent clustering patterns using the correlation features.

In this work, we focused on seizure detection and prediction

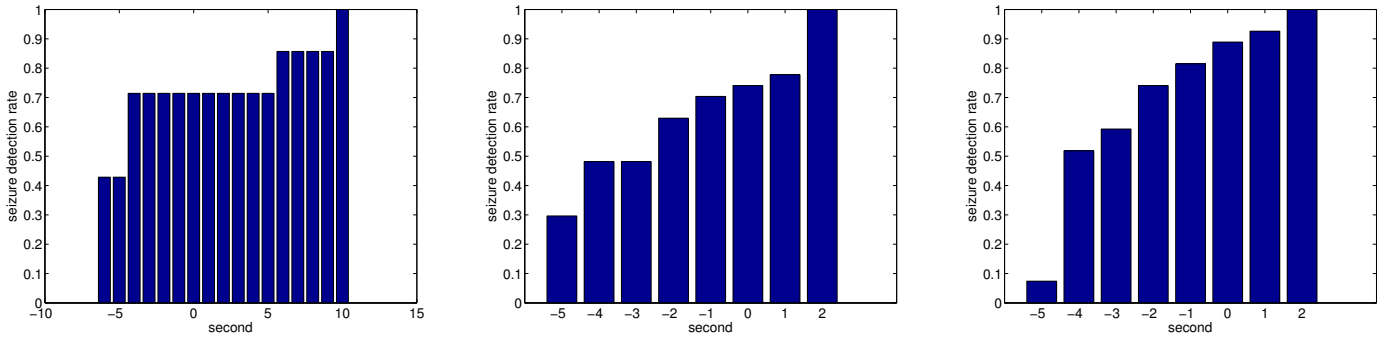

(a) dataset 1 (b) dataset 2 (c) dataset 3

Fig. 5: Histogram of seizure detection delay. Horizontal axis indicate the detection delay. Negative delay means that the seizure is predicted before the actual seizure onset. Vertical axis is the cumulative distribution of the delay. For datasets 2 and 3, all seizures were detected with up to 2 seconds delay, and more than 70% of seizures were predicted before the actual onsets. However, for dataset 1, 5 out of 7 seizures were predicted before the seizure onset, and all seizures were detected within 10 sec. delay.

Dataset	prediction accuracy	number of false positive	average time before seizure
cat 1	5/7: 0.7143	0	1.86 sec
cat 2	20/27: 0.7407	1	2.22 sec
cat 2	24/27: 0.8889	0	2.33 sec

TABLE IV: Seizure prediction accuracy for three datasets using the HMM+SVM approach. Seizure onset is predicted if 5 consecutive spikes are classified to either pre-seizure or seizure stage. The time between this prediction and the actual onset time is defined as the delay. Seizure prediction is considered correct if -8 sec $\leq$ delay $\leq$ 0 sec. The last column reports the average of the negative delays among all predictions (including those with positive delays).

using each spike segment's cluster label. We found that a two stage HMM+SVM classifier had a much higher spike classification accuracy than a single HMM classifier. The two stage approach yielded accurate seizure detection and possibly prediction, by detecting 49 out of 61 seizures before the expert-labeled seizure onset. However, since our seizure prediction was less than 5 seconds before the expert-labeled seizure onset, our seizure prediction might reflect human error and uncertainty in identifying the start of the seizure, rather than true prediction. Our short seizure prediction time horizon could also be caused by the absence of a significant pre-seizure state in our dataset. The seizures in our dataset are artificially induced using picrotoxin, causing severe abnormalities in the inter-ictal μECoG waveforms. The lack of normal background neural signals between seizures could make seizure prediction in this model difficult or impossible. By combining the methods we have developed with data from more accurate animal models of epilepsy with spontaneous seizures, our framework could yield improved seizure prediction accuracy.

V. ACKNOWLEDGEMENT

This work was funded by National Science Foundation award CCF-1422914.

REFERENCES

[1] J. Viventi, D.-H. Kim, L. Vigeland, E. S. Frechette, J. A. Blanco, Y.-S. Kim, A. E. Avrin, V. R. Tiruvadi, S.-W. Hwang, A. C. Vanleer *et al.*, "Flexible, foldable, actively multiplexed, high-density electrode array for mapping brain activity in vivo," *Nature neuroscience*, vol. 14, no. 12, pp. 1599–1605, 2011.

[2] B. Akyildiz, Y. Song, J. Viventi, and Y. Wang, "Improved clustering of spike patterns through video segmentation and motion analysis of micro electrocorticographic data," in *Signal Processing in Medicine and Biology Symposium (SPMB), 2013 IEEE*. IEEE, 2013, pp. 1–6.

[3] R. Adams and L. Bischof, "Seeded region growing," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 16, no. 6, pp. 641–647, 1994.

[4] H. Adeli, Z. Zhou, and N. Dadmehr, "Analysis of eeg records in an epileptic patient using wavelet transform," *Journal of neuroscience methods*, vol. 123, no. 1, pp. 69–87, 2003.

[5] P. J. Durka, "From wavelets to adaptive approximations: time-frequency parametrization of eeg," *BioMedical Engineering OnLine*, vol. 2, no. 1, p. 1, 2003.

[6] S. Ghosh-Dastidar, H. Adeli, and N. Dadmehr, "Mixed-band wavelet-chaos-neural network methodology for epilepsy and epileptic seizure detection," *Biomedical Engineering, IEEE Transactions on*, vol. 54, no. 9, pp. 1545–1551, 2007.

[7] Y. Park, L. Luo, K. K. Parhi, and T. Netoff, "Seizure prediction with spectral power of eeg using cost-sensitive support vector machines," *Epilepsia*, vol. 52, no. 10, pp. 1761–1770, 2011.

[8] R. J. Martis, U. R. Acharya, J. H. Tan, A. Petznick, L. Tong, C. K. Chua, and E. Y. K. Ng, "Application of intrinsic time-scale decomposition (itd) to eeg signals for automated seizure prediction," *International journal of neural systems*, vol. 23, no. 05, p. 1350023, 2013.

[9] K. Gadhoumi, J. Gotman, and J. M. Lina, "Scale invariance properties of intracerebral eeg improve seizure prediction in mesial temporal lobe epilepsy," *PloS one*, vol. 10, no. 4, 2015.

[10] M. Arnold, X. Milner, H. Witte, R. Bauer, and C. Braun, "Adaptive ar modeling of nonstationary time series by means of kalman filtering," *Biomedical Engineering, IEEE Transactions on*, vol. 45, no. 5, pp. 553–562, 1998.

[11] L. Chisci, A. Mavino, G. Perferi, M. Sciandrone, C. Anile, G. Colicchio, and F. Fuggetta, "Real-time epileptic seizure prediction using ar models and support vector machines," *Biomedical Engineering, IEEE Transactions on*, vol. 57, no. 5, pp. 1124–1132, 2010.

[12] J. R. Williamson, D. W. Bliss, and D. W. Browne, "Epileptic seizure prediction using the spatiotemporal correlation structure of intracranial eeg," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 665–668.

[13] J. R. Williamson, D. W. Bliss, D. W. Browne, and J. T. Narayanan, "Seizure prediction using eeg spatiotemporal correlation structure," *Epilepsy & Behavior*, vol. 25, no. 2, pp. 230–238, 2012.

[14] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
[15] A. W. Chiu, M. Derchansky, M. Cotic, P. L. Carlen, S. O. Turner, and B. L. Bardakjian, "Wavelet-based gaussian-mixture hidden markov model for the detection of multistage seizure dynamics: A proof-of-concept study," *Biomed Eng Online*, vol. 10, p. 29, 2011.
[16] S. Wong, A. B. Gardner, A. M. Krieger, and B. Litt, "A stochastic framework for evaluating seizure prediction algorithms using hidden markov models," *Journal of neurophysiology*, vol. 97, no. 3, pp. 2525–2532, 2007.
[17] S. Santaniello, S. P. Burns, A. J. Golby, J. M. Singer, W. S. Anderson, and S. V. Sarma, "Quickest detection of drug-resistant seizures: An optimal control approach," *Epilepsy & Behavior*, vol. 22, pp. S49–S60, 2011.
[18] S. Santaniello, D. L. Sherman, N. V. Thakor, E. N. Eskandar, and S. V. Sarma, "Optimal control-based bayesian detection of clinical and behavioral state transitions," *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 20, no. 5, pp. 708–719, 2012.
[19] D. M. Blei, M. I. Jordan *et al.*, "Variational inference for dirichlet process mixtures," *Bayesian analysis*, vol. 1, no. 1, pp. 121–143, 2006.